

From FITS to SQL - Loading and Publishing the SDSS Data

Aniruddha R. Thakar and Alexander S. Szalay

*Center for Astrophysical Sciences, The Johns Hopkins University,
Baltimore, MD 21218, Email: thakar@pha.jhu.edu*

Jim Gray

Microsoft Research

Abstract. For large astronomical databases like the SDSS Science Archive, data loading is potentially the most time-consuming and labor-intensive part of archive operations, and it is also the most critical: it is the last chance to examine the data before it is published. We attempted to automate this job as much as possible, and to make it easy to diagnose data and loading errors. We describe the **sqlLoader** - a distributed workflow system of modules that check, load, validate and publish the data to the databases. The workflow is described by a directed acyclic graph whose nodes are the processing modules. It is designed for parallel loading and is controlled from a web interface (Load Monitor). The validation stage represents a systematic and thorough scrubbing of the data. Finally, the different data products are merged into a set of linked tables that can be efficiently searched with specialized indices and pre-computed joins.

1. Introduction

The Sloan Digital Sky Survey Data Release 1 (DR1) contains a Terabyte of catalog data published online at <http://skyserver.sdss.org/dr1>. The total size of the catalog data including backup and archive copies is about 3-4 TB, and presents special challenges in terms of its storage and organization (Thakar et al. 2003) as well as loading and publishing it. The load/publish process is very time-consuming for Terabyte databases (several days to several weeks at current disk speeds), and data once publicly distributed cannot be taken offline or changed after it has been used for scientific publications.

Loading the data is therefore a critical step whose efficiency and success are paramount to the timely availability and correctness of the published data. However, in reality it is often the least automated, least budgeted and planned-for step in archive operations, and almost always takes far longer than anticipated.

Unfortunately, loading is not a one-time thing. Data errors cause reloads. Errors in the software pipeline cause reloads. Sometimes, the fastest way to redesign and reorganize the database is to simply reload it. So it is very important to have a fully automated load process that can reload the entire database in a few days with minimal supervision. Accordingly, we have developed an au-

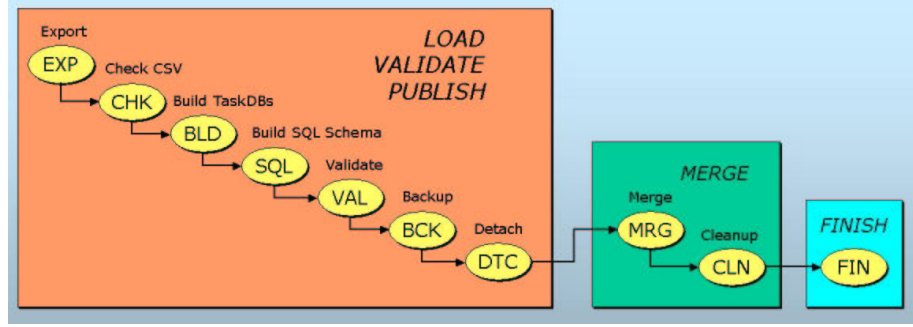


Figure 1. sqlLoader steps represented as a Directed Acyclic Graph.

tomated loading and publishing pipeline called **sqlLoader** that takes the data exported in FITS format by the SDSS Operational DB, converts it to CSV (comma-separated values) format, then validates and publishes it to the Microsoft SQL Server DBMS that contains our catalog data. The heart of the sqlLoader is a distributed workflow system of modules described by a directed acyclic graph (DAG) (Figure 1).

2. Loading Process

The basic processing entity is a *task*. A task is started when a data *chunk* is exported by the OpDB. Exported chunks are converted to CSV format, and are contained in a single directory. There are several different export types - TARGET, BEST, RUNS, PLATE and TILING - each of which is subjected to different tests and destined for a different database. Each task comes with a taskID that is unique within its category.

The loading process consists of *steps*. The first step is to **load** each chunk of data into a separate task DB, containing only a thin set of indices. Then we **validate** the data. This includes verifying that there are no primary key collisions and all foreign keys point to a valid record. As we find problems downstream we add more and more tests for data ranges, sanity checks, and cross-checks on cardinalities. We build several ancillary tables for spatial searches (HTM, Neighbors, etc.) After the validation step we **publish** the data: we perform a DB-to-DB copy, where the target is the final production database. After publishing, we make a **backup** of the task DB. At the very end, all the different datasets (databases) are **merged** into the final (published) database in the finish step, and indices are created for efficient data mining.

3. Validation

Validation is perhaps the most important step in the loading process. The speed, integrity and convenience that databases offer come at a price: data once published cannot be retracted or changed. The data must always be available once science has been done with it, and it must be correct because publications are based on it. The validate step in sqlLoader represents a systematic scrubbing and sanity-check of the data, from a scientific as well as data integrity point

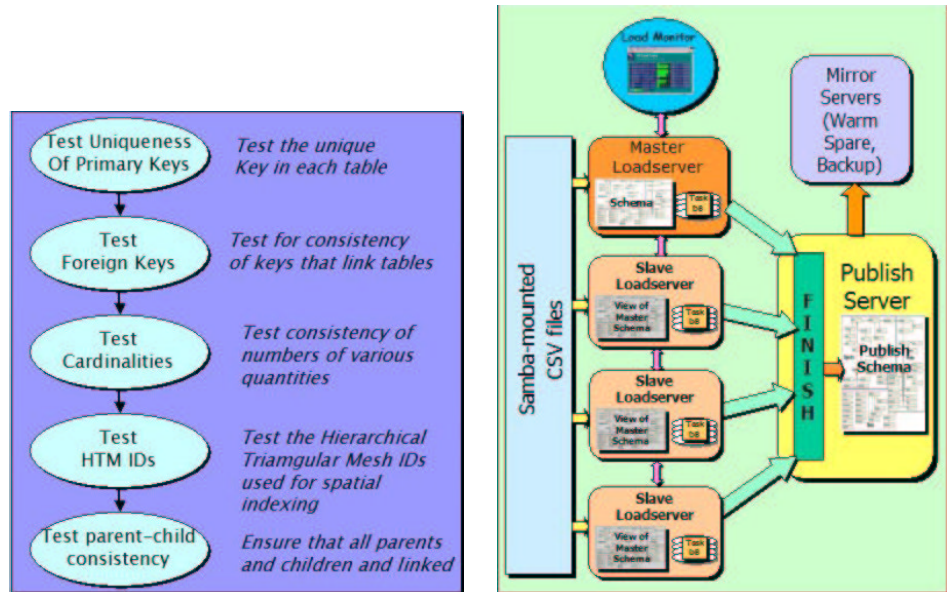


Figure 2. (a) Data validation steps and (b) Distributed Loading.

of view. Figure 2(a) shows the various operations that are performed on the data. The primary and foreign key tests are run on all the tables. The photo (imaging) and spectro tables are tested for HTM IDs, which are 64-bit IDs used by the Hierarchical Triangular Mesh indexing scheme. The referential integrity of parent-child relationships for debleded image objects is also checked. Finally, the consistency of counts of various quantities is checked.

This validation process has proven invaluable in finding numerous inconsistencies and errors in the data and catching them early, during the testing of DR1 rather than after the data is published. If a chunk fails somewhere in the load process, the operator can kill the operation and restart it. Every step of the pipeline generates a 3-level detailed log that records the operation, how long it took and how many errors were detected. This allows the operators to quickly identify data quality problems, performance problems, and “stuck” tasks.

4. Distributed Loading

Loading a Terabyte or more of data is a time-consuming process even with fast disks, and parallelization of the loading steps is a big help, especially as we get into the multi-TB data volumes of future SDSS releases. The load, validate and publish steps in the sqlLoader (leftmost rectangle in Fig. 1) are fully parallelizable and can be executed on a distributed cluster of load-servers. Figure 2(b) shows the setup for parallel distributed loading. The master schema is on the loadadmin (master) server, with each loadsupport (slave) having a remote *view* of the schema. Distributed loading makes use of advanced DBMS features like linked servers with distributed views and distributed transactions. After loading, validating and publishing is done in parallel, the merging of the parallel data streams and the finish step are performed sequentially on the publish server.

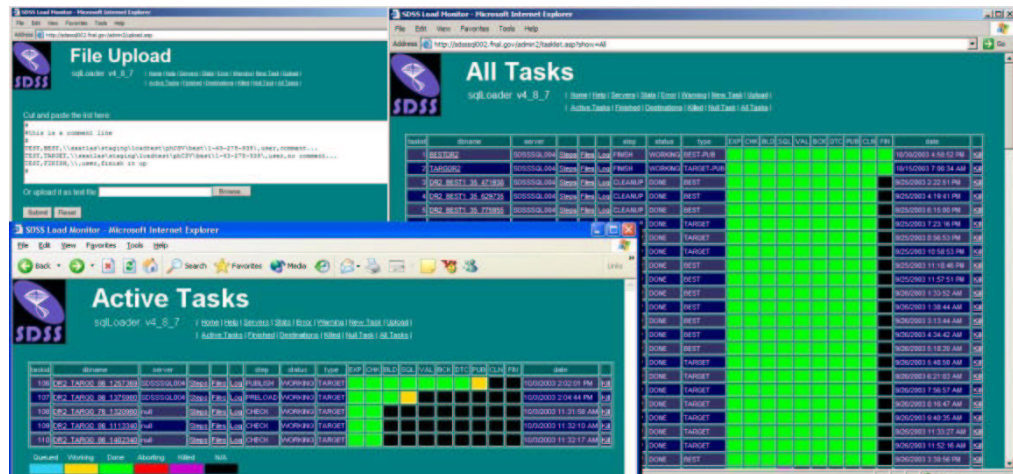


Figure 3. Loader Monitor screens showing All Tasks, Active Tasks and File Upload pages.

5. The Load Monitor

The Load Monitor is the admin web interface to the sqlLoader. It enables job submission, control and tracking via a user-friendly GUI. Loading jobs (tasks) can be submitted either a single chunk at a time or bulk-uploaded with the File Upload feature. Tasks can be monitored at several levels, and information is available on the status of the individual files being loaded, detailed logs of each step in the task, and separate listings of errors and warnings encountered. Sample Load Monitor screens are shown in Figure 3 to illustrate the features that are available. The Active Tasks listing shows all the tasks that are currently executing. The task status is displayed as a color code for each step (export, check, build, validate, backup, detach, merge, cleanup and finish). Amber means the step is in progress, green means it is done, red means it failed and purple means that the task was killed in that step. The Servers page shows status of each server and enables individual servers to be started and stopped.

6. Conclusions

The sqlLoader pipeline is a distributed workflow system that automates the loading and publishing of the SDSS data. It has enabled the loading for DR1 to be completed largely as a turnkey operation with very little human intervention.

References

- Thakar, A.R., Szalay, A.S., vandenBerg, J.V. & Gray, J. 2003, in APS Conf. Ser., Vol. 295, Astronomical Data Analysis Software and Systems XII, ed. H. E. Payne, R. I. Jedrzejewski, & R. N. Hook (San Francisco: ASP), 217.